

Gaussian-Dirichlet Posterior Dominance in Sequential Learning

Ian Osband*
Stanford University
ian.osband@gmail.com

Benjamin Van Roy
Stanford University
bvr@stanford.edu

February 15, 2017

Abstract

We consider the problem of sequential learning from categorical observations bounded in $[0, 1]$. We establish an ordering between the Dirichlet posterior over categorical outcomes and a Gaussian posterior under observations with $N(0, 1)$ noise. We establish that, conditioned upon identical data with at least two observations, the posterior mean of the categorical distribution will always second-order stochastically dominate the posterior mean of the Gaussian distribution. These results provide a useful tool for the analysis of sequential learning under categorical outcomes.

1 Introduction

For any $S \in \mathbb{N}$, fix any $v \in [0, 1]^S$ and consider probabilities P_1, \dots, P_S associated with components of v . Let the vector P of probabilities itself be random and Dirichlet-distributed with parameters $\alpha \in \mathbb{R}_{++}^S$. Let $X|P$ be a vector of n independent samples drawn from the associated categorical distribution over components of v . Note that the components of X are conditionally independent, conditioned on P , but are not unconditionally independent. Conditioned on P , the mean of each X_n is $\bar{X} = P^\top v$. Let $z \in [0, 1]^n$ and $c_s = \sum_{i=1}^n \mathbb{1}\{z_i = v_s\}$ for each $s = 1, \dots, S$. Then, the distribution of P conditioned on $X = z$ is Dirichlet with parameters $\alpha + c$.

Let \bar{Y} be distributed $N(\mathbb{E}[\bar{X}], (\alpha^\top \mathbf{1})^{-1})$. Let $Y|\bar{Y}$ be a vector of N independent samples distributed according to $N(\bar{Y}, 1)$. The distribution of \bar{Y} conditioned on $Y = z$ is $N(\mu, \sigma^2)$, where

$$\mu = \frac{(\alpha + c)^\top v}{(\alpha + c)^\top \mathbf{1}} \quad \text{and} \quad \sigma^2 = \frac{1}{(\alpha + c)^\top \mathbf{1}}.$$

In this paper, we establish that $\bar{X}|(X=z) \succ_{\text{ssd}} \bar{Y}|(Y=z)$, where \succ_{ssd} denotes second-order stochastic dominance. In other words, conditioned on identical outcomes, the posterior mean of the categorical distribution second-order-stochastically dominates the posterior mean of the Gaussian distribution.

*Corresponding author, now at DeepMind.

This result extends earlier work relating variances of posterior means under Gaussian and Dirichlet models (Antoniak, 1974; Kyung et al., 2009). Our result provides a dominance relation that applies to all moments. Our interest in this result stems from its significance in the area of *reinforcement learning* (Sutton and Barto, 1998), where we have used it to establish a notion of stochastic optimism achieved by particular reinforcement algorithms that generate randomized value functions to explore in an efficient manner (Osband et al., 2014; Osband, 2016). This paper presents the result and its proof in a form that will be cited by our work on reinforcement learning and that will be accessible to researchers more broadly.

2 Stochastic dominance

In this section we will review several notions of partial orderings for real-valued random variables. All random variables we define will be with respect to the probability space $(\Omega, \mathcal{F}, \mathbb{P})$.

Definition 1 (First order stochastic dominance (FSD)).

Let X and Y be real-valued random variables. We say that X is (first order) stochastically dominant for Y if for all $x \in \mathbb{R}$,

$$\mathbb{P}(X > x) \geq \mathbb{P}(Y > x). \quad (1)$$

We write $X \succsim_{\text{fsd}} Y$ for this relationship.

First order stochastic dominance defines a partial ordering between random variables but it also quite a blunt notion of dominance that will be insufficient for our purposes. Consider $X \sim N(0, \sigma_X^2)$ and $Y \sim N(0, \sigma_Y^2)$ with $\sigma_X < \sigma_Y$. These random variables cannot be related in terms of FSD. However, in the context of gambling we might imagine that the return from X is in some sense preferable to Y , since they have the same mean but X is somehow less risky. Our next definition formalizes this notion.

Definition 2 (Second order stochastic dominance (SSD)).

Let X and Y be real-valued random variables. We say that X is second order stochastically dominant for Y if for all $u : \mathbb{R} \rightarrow \mathbb{R}$ convex and non-decreasing,

$$\mathbb{E}[u(X)] \geq \mathbb{E}[u(Y)]. \quad (2)$$

We write $X \succsim_{\text{ssd}} Y$ for this relationship.

Proposition 1 (SSD equivalence).

Let X and Y be real-valued random variables with finite expectation. The following are equivalent:

1. $X \succsim_{\text{ssd}} Y$
2. For any $u : \mathbb{R} \rightarrow \mathbb{R}$ concave and increasing $\mathbb{E}[u(X)] \geq \mathbb{E}[u(Y)]$
3. For any $\alpha \in \mathbb{R}$, $\int_{-\infty}^{\alpha} \{\mathbb{P}(Y \leq s) - \mathbb{P}(X \leq s)\} ds \geq 0$.
4. $Y =_D X + A + W$ for $A \leq 0$ and $\mathbb{E}[W|X + A] = 0$ for all values $x + a$.

Proof. This follows from a simple integration by parts (Hadad and Russell, 1969). □

Second order stochastic dominance $X \succsim_{\text{ssd}} Y$ ensures that $\mathbb{E}[X] \geq \mathbb{E}[Y]$. It also establishes that for any convex loss L that X is less “spread out” than Y in the sense $\mathbb{E}[L(X - \mathbb{E}[X])] \leq \mathbb{E}[L(Y - \mathbb{E}[Y])]$. Motivated by this equivalence, we introduce another related dominance condition.

Definition 3 (Single crossing dominance (SCD)).

Let X and Y be real-valued random variables with CDFs F_X, F_Y and finite expectation. We say that X single-crossing dominates Y if and only if $\mathbb{E}[X] \geq \mathbb{E}[Y]$ and there a crossing point $a \in \mathbb{R}$ such that:

$$F_Y(s) \geq F_X(s) \iff s \leq a. \quad (3)$$

We write $X \succ_{sc} Y$ for this relationship.

Single crossing dominance is actually a stronger condition than SSD, as we show in Proposition 2. In general the reverse implication is not true, as we demonstrate in Example 1.

Proposition 2 (SCD implies SSD).

Let X and Y be real-valued random variables with finite expectation then

$$X \succ_{sc} Y \implies X \succ_{ssd} Y. \quad (4)$$

Proof. Suppose $X \succ_{sc} Y$ with single crossing point a . Let $I(\alpha) = \int_{-\infty}^{\alpha} \{\mathbb{P}(Y \leq s) - \mathbb{P}(X \leq s)\} ds$. By $X \succ_{sc} Y$ we know $I(\alpha) > 0$ for all $\alpha \leq a$ and that $I(\alpha)$ is decreasing for all $\alpha \geq a$. Now we consider the limit $\lim_{\alpha \rightarrow \infty} I(\alpha) = \mathbb{E}[X] - \mathbb{E}[Y] \geq 0$. Hence $I(\alpha) \geq 0$ for all $\alpha \in \mathbb{R}$, which shows that $X \succ_{so} Y$ by Proposition 1. \square

Example 1 (SSD does not imply SCD).

Consider $X \sim \text{Unif}(\{-1, 1\})$ and let $Y = X + W$ where $W \sim \text{Unif}([-1, 1])$ and independent of Y . By Proposition 1 $X \succ_{ssd} Y$, however X is not single crossing dominant for Y .

Proof. We display the CDFs of these variables in Figure 1, they are not single crossing. In particular the ordering of F_X, F_Y switches at least three points $x = -1, 0, 1$. \square




Figure 1: Second order stochastic dominance does not imply single crossing dominance

3 Gaussian-Dirichlet dominance

The main technical result in this paper comes in Theorem 1, which we prove in Section 4.

Theorem 1 (Gaussian vs Dirichlet dominance).

Let $X = P^T v$ for $v \in [0, 1]^S$ fixed and $P \sim \text{Dirichlet}(\alpha)$ with $\alpha \in \mathbb{R}_+^S$ and $\sum_{i=1}^S \alpha_i \geq 2$. Let $Y \sim N(\mu, \sigma^2)$ with $\mu = \frac{\sum_{i=1}^S \alpha_i v_i}{\sum_{i=1}^S \alpha_i}$, $\sigma^2 = \left(\sum_{i=1}^S \alpha_i\right)^{-1}$, then $\mathbb{E}[X] = \mathbb{E}[Y]$ and $X \succ_{\text{ssd}} Y$.

At first glance, Theorem 1 may seem quite arcane, it provides an ordering between two paired families of Gaussian and Dirichlet distributions in terms of SSD. The reason this result is so useful is that, given matched prior distributions, the resultant posteriors for the Gaussian and Dirichlet models will remain ordered in this way for any observation data. The condition $\sum_{i=1}^S \alpha_i \geq 2$ is technical but does not pose significant difficulties so long as at the posterior is updated with at least two observations. We present this result as Corollary 1.

Corollary 1 (Gaussian vs Dirichlet posterior ordering).

Let $X = P^T v$ for $v \in [0, 1]^S$ fixed and $P \sim \text{Dirichlet}(\alpha)$ with $\alpha \in \mathbb{R}_+^S$. Let $Y \sim N(\mu, \sigma^2)$ with $\mu = \frac{\sum_{i=1}^S \alpha_i v_i}{\sum_{i=1}^S \alpha_i}$, $\sigma^2 = \left(\sum_{i=1}^S \alpha_i\right)^{-1}$. Let D be the data from n i.i.d. samples from the categorical distribution P and values v . Let \tilde{X} be the posterior distribution for $X \mid D$ and \tilde{Y} be the posterior distribution for $Y \mid D$ but updating according to a mis-specified likelihood as if the observations were $\sim N(\mu, 1)$. Then, for all datasets D such that $n + \sum_{i=1}^S \alpha_i \geq 2$ we can guarantee that $\tilde{X} \succ_{\text{ssd}} \tilde{Y}$.

Proof. This result is a consequence of Theorem 1 together with algebraic relations for the conjugate updates of \tilde{X} and \tilde{Y} given any data D . We write $\vec{n} \in \mathbb{N}^S$ for the number of observations from each category v in the dataset D with $\sum_{s=1}^S \vec{n}_s = n$. Then we can write the posterior distribution $\tilde{X} = X \mid D = P^T v$ for $P \sim \text{Dirichlet}(\tilde{\alpha})$ and $\tilde{\alpha} := \alpha + \vec{n}$.

In a similar way we can compute the posterior distribution of $\tilde{Y} = Y \mid D$ where we update with an misspecified likelihood as if the data were $\sim N(\mu, 1)$. Once again we can use a conjugate form for the update $\tilde{Y} \sim N(\mu, \sigma^2)$ explicitly,

$$\mu = \frac{\sum_{s=1}^S \vec{n}_s v_s + \frac{\sum_{s=1}^S \alpha_s v_s}{\sum_{s=1}^S \alpha_s} \sum_{s=1}^S \alpha_s}{\sum_{s=1}^S \vec{n}_s + \sum_{s=1}^S \alpha_s} = \frac{\sum_{s=1}^S \tilde{\alpha}_s v_s}{\sum_{s=1}^S \tilde{\alpha}_s},$$

$$\sigma^2 = \frac{1}{\sum_{s=1}^S \vec{n}_s + \sum_{s=1}^S \alpha_s} = \frac{1}{\sum_{s=1}^S \tilde{\alpha}_s}.$$

We conclude by application of Theorem 1 on the updated posterior parameters $\tilde{\alpha}$. □

4 Proof of Theorem 1

The complete proof of Theorem 1 is long but the essential argument is simple. We outline the main arguments below and fill in the details in Sections 4.1 and 4.2. First, we consider an auxilliary random variable $\tilde{X} \sim \text{Beta}(\tilde{\alpha}, \tilde{\beta})$ with $\tilde{\alpha} = \sum_{i=1}^S \alpha_i v_i$ and $\tilde{\beta} = \sum_{i=1}^S \alpha_i - \tilde{\alpha}$. In Lemma 2 we show

that $X \succ_{\text{ssd}} \tilde{X}$. Next, we show that this auxiliary beta \tilde{X} is single crossing dominant for the approximating Gaussian posterior, $\tilde{X} \succ_{\text{sc}} Y$. Therefore, by Proposition 2 $X \succ_{\text{ssd}} Y$.

The main difficulty in this proof comes in establishing $\tilde{X} \succ_{\text{sc}} Y$. To do this we use a laborious calculus argument together with repeated applications of the mean value theorem. Our proof requires separate upper and lower bounds for different regions of $\tilde{\alpha}$ and $\tilde{\beta}$, but no real insight beyond that. We believe that there should be a much more enlightened and elegant method to obtain these results.

4.1 Beta vs Dirichlet

We begin our proof of Theorem 1 with an intermediate comparison of the Dirichlet distribution to a matched Beta posterior. We first state a more basic result that we will use on Gamma distributions.

Lemma 1 (Conditioning the sum of Gamma random variables).

Let $\gamma_1 \sim \text{Gamma}(k_1, \theta)$ and $\gamma_2 \sim \text{Gamma}(k_2, \theta)$ be independent random variables. Then the conditional expectations $\mathbb{E}[\gamma_1 | \gamma_1 + \gamma_2] = \frac{k_1}{k_1 + k_2}(\gamma_1 + \gamma_2)$ and $\mathbb{E}[\gamma_2 | \gamma_1 + \gamma_2] = \frac{k_2}{k_1 + k_2}(\gamma_1 + \gamma_2)$.

Lemma 2 (Beta vs Dirichlet dominance).

Let $X = P^\top v$ for the random variable $P \sim \text{Dirichlet}(\alpha)$ and constants $v \in \mathbb{R}^S$ and $\alpha \in \mathbb{R}_+^S$. Without loss of generality, assume $v_1 \leq v_2 \leq \dots \leq v_S$. Let $\tilde{\alpha} = \sum_{i=1}^S \alpha_i(v_i - v_1)/(v_d - v_1)$ and $\tilde{\beta} = \sum_{i=1}^d \alpha_i(v_d - v_i)/(v_d - v_1)$. Then, there exists a random variable $\tilde{p} \sim \text{Beta}(\tilde{\alpha}, \tilde{\beta})$ such that, for $\tilde{X} = \tilde{p}v_d + (1 - \tilde{p})v_1$, $\mathbb{E}[\tilde{X}|X] = \mathbb{E}[X]$ and so $X \succ_{\text{ssd}} \tilde{X}$.

Proof. Let $\gamma_i \sim \text{Gamma}(\alpha_i, 1)$, with $\gamma_1, \dots, \gamma_d$ independent, and let $\bar{\gamma} = \sum_{i=1}^d \gamma_i$, so that $P \equiv_D \gamma/\bar{\gamma}$. Let $\alpha_i^0 = \alpha_i(v_i - v_1)/(v_d - v_1)$ and $\alpha_i^1 = \alpha_i(v_d - v_i)/(v_d - v_1)$ so that $\alpha = \alpha^0 + \alpha^1$. Define independent random variables $\gamma^0 \sim \text{Gamma}(\alpha_i^0, 1)$ and $\gamma^1 \sim \text{Gamma}(\alpha_i^1, 1)$ so that $\gamma \equiv_D \gamma^0 + \gamma^1$.

Take γ^0 and γ^1 to be independent, and couple these variables with γ so that $\gamma = \gamma^0 + \gamma^1$. Note that $\tilde{\beta} = \sum_{i=1}^d \alpha_i^0$ and $\tilde{\alpha} = \sum_{i=1}^d \alpha_i^1$. Let $\bar{\gamma}^0 = \sum_{i=1}^d \gamma_i^0$ and $\bar{\gamma}^1 = \sum_{i=1}^d \gamma_i^1$, so that $1 - \tilde{p} \equiv_D \bar{\gamma}^0/\bar{\gamma}$ and $\tilde{p} \equiv_D \bar{\gamma}^1/\bar{\gamma}$. Couple these variables so that $1 - \tilde{p} = \bar{\gamma}^0/\bar{\gamma}$ and $\tilde{p} = \bar{\gamma}^1/\bar{\gamma}$. We can now say

$$\begin{aligned} \mathbb{E}[\tilde{X}|X] &= \mathbb{E}[(1 - \tilde{p})v_1 + \tilde{p}v_d|X] = \mathbb{E}\left[\frac{v_1\bar{\gamma}^0}{\bar{\gamma}} + \frac{v_d\bar{\gamma}^1}{\bar{\gamma}} \middle| X\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\frac{v_1\bar{\gamma}^0 + v_d\bar{\gamma}^1}{\bar{\gamma}} \middle| \gamma, X\right] \middle| X\right] = \mathbb{E}\left[\frac{v_1\mathbb{E}[\bar{\gamma}^0|\gamma] + v_d\mathbb{E}[\bar{\gamma}^1|\gamma]}{\bar{\gamma}} \middle| X\right] \\ &= \mathbb{E}\left[\frac{v_1 \sum_{i=1}^d \mathbb{E}[\gamma_i^0|\gamma_i] + v_d \sum_{i=1}^d \mathbb{E}[\gamma_i^1|\gamma_i]}{\bar{\gamma}} \middle| X\right] \\ &\stackrel{(a)}{=} \mathbb{E}\left[\frac{v_1 \sum_{i=1}^d \gamma_i \alpha_i^0 / \alpha_i + v_d \sum_{i=1}^d \gamma_i \alpha_i^1 / \alpha_i}{\bar{\gamma}} \middle| X\right] \\ &= \mathbb{E}\left[\frac{v_1 \sum_{i=1}^d \gamma_i (v_i - v_1) + v_d \sum_{i=1}^d \gamma_i (v_d - v_i)}{\bar{\gamma}(v_d - v_1)} \middle| X\right] \\ &= \mathbb{E}\left[\frac{\sum_{i=1}^d \gamma_i v_i}{\bar{\gamma}} \middle| X\right] = \mathbb{E}\left[\sum_{i=1}^d p_i v_i \middle| X\right] = X, \end{aligned}$$

where (a) follows from Lemma 1. Therefore, \tilde{X} is a mean-preserving spread of X and so by Proposition 1, $X \succ_{\text{ssd}} \tilde{X}$. \square

4.2 Gaussian vs Beta

We complete the proof of Theorem 1 by showing that this auxiliary Beta random variable \tilde{X} defined in Lemma 2 is second order stochastic dominant for the Gaussian posterior Y .

Lemma 3 (Gaussian vs Beta dominance).

Let $\tilde{X} \sim \text{Beta}(\alpha, \beta)$ for any $\alpha > 0, \beta > 0$ and $Y \sim N\left(\mu = \frac{\alpha}{\alpha+\beta}, \sigma^2 = \frac{1}{\alpha+\beta}\right)$. Then, $\tilde{X} \succ_{\text{sc}} Y$ (and by Proposition 2 this implies $\tilde{X} \succ_{\text{ssd}} Y$) whenever $\alpha + \beta \geq 2$.

We want to prove that the CDFs cross at most once on $(0, 1)$. By the mean value theorem (Rudin, 1964), it is sufficient to prove that the PDFs cross at most twice on the same interval. We lament that the proof as it stands is so laborious, but our attempts at a more elegant solution has so far been unsuccessful. The remainder of this appendix is devoted to proving this “double-crossing” property via manipulation of the PDFs for different values of α, β .

We write f_N for the density of the Normal Y and f_B for the density of the Beta \tilde{X} respectively. We know that at the boundary $f_N(0-) > f_B(0-)$ and $f_N(1+) > f_B(1+)$ where the \pm represents the left and right limits respectively. As these densities are positive over the interval, we can consider the log PDFs

$$\begin{aligned} l_B(x) &= (\alpha - 1)\log(x) + (\beta - 1)\log(1 - x) + K_B \\ l_N(x) &= -\frac{1}{2}(\alpha + \beta) \left(x - \frac{\alpha}{\alpha + \beta}\right)^2 + K_N. \end{aligned}$$

The function $\log(x)$ is injective and increasing; if we can show that $l_N(x) - l_B(x) = 0$ has at most two solutions on the interval we will be done.

Instead we will attempt to prove an even stronger condition, that $l'_N(x) - l'_B(x) = 0$ has at most one solution in the interval. This sufficient condition may be easier to deal with since we can ignore the distributional normalizing constants.

$$l'_B(x) = \frac{\alpha - 1}{x} - \frac{\beta - 1}{1 - x}, \quad l'_N(x) = \alpha - (\alpha + \beta)x$$

Finally we consider an even stronger condition, if $l''_N(x) - l''_B(x) = 0$ has no solution then $l'_B(x) - l'_N(x)$ must be monotone over the region and so it can have at most one root.

$$l''_B(x) = -\frac{\alpha - 1}{x^2} - \frac{\beta - 1}{(1 - x)^2}, \quad l''_N(x) = -(\alpha + \beta)$$

With these definitions now let us define:

$$h(x) := l''_N(x) - l''_B(x) = \frac{\alpha - 1}{x^2} + \frac{\beta - 1}{(1 - x)^2} - (\alpha + \beta) \quad (5)$$

Our goal now is to show that $h(x) = 0$ does not have any solutions for $x \in [0, 1]$. Once again, we will look at the derivatives and analyze them for different values of $\alpha, \beta > 0$.

$$\begin{aligned} h'(x) &= -2 \left(\frac{\alpha - 1}{x^3} - \frac{\beta - 1}{(1 - x)^3} \right) \\ h''(x) &= 6 \left(\frac{\alpha - 1}{x^4} + \frac{\beta - 1}{(1 - x)^4} \right) \end{aligned}$$

Our proof will proceed by considering specific ranges for the values of $\alpha, \beta > 0$ and use different calculus arguments for each of these regions. By symmetry in the problem, we only need to prove the result for $\alpha > \beta$. Within this section of possible parameter values we will need to subdivide the quadrant into three proof regions. $R1 := \{\alpha > 1 \geq \beta \geq 0\}$, $R2 := \{\alpha > 1, \beta > 1, (\alpha - 1)(\beta - 1) \geq 1/9\}$ and $R3 := \{\alpha > 1, \beta > 1, (\alpha - 1)(\beta - 1) < 1/9\}$. These regions completely cover all $\alpha + \beta \geq 2$ and hence suffice to complete the proof of Lemma 3.




Figure 2: Parameter regions for proof. The special case $\alpha = \beta = 1$ can be verified individually, in this case the PDFs do not intersect at any point.

4.2.1 Region $R1 = \{\alpha \geq 1 \geq \beta \geq 0\}$

In this region we will show that $g(x) = l'_N(x) - l'_B(x)$ has no solutions. We write $A = \alpha - 1 > 0$ and $B = \beta - 1 \leq 0$ as before.

$$g(x) = \alpha - (\alpha + \beta)x + \frac{\beta - 1}{1 - x} - \frac{\alpha - 1}{x}$$

$$g'(x) = h(x) = \frac{A}{x^2} + \frac{B}{(1 - x)^2} - (\alpha + \beta)$$

$$g''(x) = h'(x) = -2 \left(\frac{A}{x^3} - \frac{B}{(1 - x)^3} \right)$$

We note that $g''(x) \leq 0$ and so $g(x)$ is a concave function. If we can show that the maximum of g lies below 0 then we know that there can be no roots. We now attempt to solve $g'(x) = 0$:

$$g'(x) = \frac{A}{x^2} + \frac{B}{(1 - x)^2} = 0 \implies -A/B = \left(\frac{x}{1 - x} \right)^2 \implies x = \frac{K}{1 + K} \in (0, 1),$$

where here we write $K = \sqrt{-A/B} > 0$. We ignore the case $B = 0$ as a trivial special case. We write $C = -B \geq 0$ and evaluate the function g at its minimum $x_K = \frac{K}{1+K}$.

$$\begin{aligned}
g(x_K) &= (A+1) - (A+B+2)\frac{K}{1+K} + B(1+K) - A\frac{1+K}{K} \\
&= -AK^2 - AK - A + BK^3 + BK^2 + BK - K^2 + K \\
&= -AK^2 - AK - A - CK^3 - CK^2 - CK - K^2 + K \\
&= -A(A/C) - A(A/C)^{1/2} - A - C(A/C)^{3/2} - C(A/C) - C(A/C)^{1/2} - A/C + (A/C)^{1/2} \\
&= -A^2C^{-1} - A^{3/2}C^{-1/2} - A - A^{3/2}C^{-1/2} - A - A^{1/2}C^{1/2} - AC^{-1} + A^{1/2}C^{1/2} \\
&= -A^2C^{-1} - 2A^{3/2}C^{-1/2} - 2A - AC^{-1} \leq 0
\end{aligned}$$

Therefore the Lemma holds for all $\alpha, \beta \in R1$

4.2.2 Region $R2 = \{\alpha > 1, \beta > 1, (\alpha - 1)(\beta - 1) \geq 1/9\}$

In the case of $\alpha, \beta > 1$ we know that $h(x)$ is a convex function on $(0, 1)$. If we solve $h'(x^*) = 0$ and $h(x^*) > 0$ then we prove our statement. We will write $A = \alpha - 1, B = \beta - 1$ for convenience.

First we solve $h'(x) = 0$ in terms of $K = (A/B)^{1/3} > 0$,

$$h'(x) = \frac{A}{x^3} - \frac{B}{(1-x)^3} = 0 \implies A/B = \left(\frac{x}{1-x}\right)^3 \implies x = \frac{K}{1+K} \in (0, 1).$$

We can now evaluate the function h at its minimum $x_K = \frac{K}{1+K}$.

$$\begin{aligned}
h(x_K) &= A\frac{(K+1)^2}{K^2} + B(K+1)^2 - (A+B+2) \\
&= A(2/K + 1/K^2) + B(K^2 + 2K) - 2 \\
&= 3(A^{2/3}B^{1/3} + A^{1/3}B^{2/3}) - 2.
\end{aligned}$$

As long as $h(x_K) > 0$ we have shown that the CDFs are single crossing. We note that for all $\alpha, \beta \in R2$

$$A, B \geq 1/3 \implies AB \geq 1/9 \implies (A^{2/3}B^{1/3} + A^{1/3}B^{2/3}) \geq 2/3.$$

This completes the proof for $R2$.

4.2.3 Region $R3 = \{\alpha > 1, \beta > 1, (\alpha - 1)(\beta - 1) < 1/9\}$

Our argument for this final region is no different than before, although it is slightly more involved. The key additional difficulty is that it in this region is not enough to only look at the derivatives of the log likelihoods; we need to use some bound on the normalizing constants to get our bounds.

In $R3$, we know that $\beta \in (1, \frac{4}{3})$ so we will make use of an upper bound to the normalizing constant of the Beta distribution, the Beta function.

$$B(\alpha, \beta) = \int_{x=0}^1 x^{\alpha-1}(1-x)^{\beta-1}dx \leq \int_{x=0}^1 x^{\alpha-1}dx = \frac{1}{\alpha} \quad (6)$$

The intuition is that, because in $R3$ the value of $\beta - 1$ is relatively small, this approximation will not be too bad. Therefore, we can explicitly bound the log likelihood of the Beta distribution:

$$l_B(x) \geq \tilde{l}_B(x) := (\alpha - 1) \log(x) + (\beta - 1) \log(1 - x) + \log(\alpha)$$

We now repeat a familiar argument based upon explicit calculus. We want to find two points $x_1 < x_2$ for which $h(x_i) = l_N''(x) - l_B''(x) > 0$. Since $\alpha, \beta > 1$ we know that h is convex and so for all $x \notin [x_1, x_2]$ then $h > 0$. We define the gap of the Beta over the maximum of the normal log likelihood,

$$\text{Gap} : l_B(x_i) - l_N(x_i) \geq f(x_i) := \tilde{l}_B(x_i) - \max_x l_N(x) > 0. \quad (7)$$

If we can show the gap is positive then it must mean there are no crossings over the region $[x_1, x_2]$. This is because \tilde{l}_B is concave and therefore totally above the maximum of l_N over the whole region $[x_1, x_2]$.

Consider any $x \in [0, x_1]$; we know from the ordering of the tails of the CDF that if there is more than one root in this segment then there must be at least three crossings. If there are three crossings, then the second derivative of their difference h must have at least one root on this region. However we know that h is convex, so if we can show that $h(x_i) > 0$ this cannot be possible. We use a similar argument for $x \in (x_2, 1]$ and complete this proof via laborious calculus.

We remind the reader of the definition in (5), $h(x) := l_N''(x) - l_B''(x) = \frac{\alpha-1}{x^2} + \frac{\beta-1}{(1-x)^2} - (\alpha + \beta)$. For ease of notation we will write $A = \alpha - 1, B = \beta - 1$. We note that:

$$h(x) > h_1(x) = \frac{A}{x^2} - (A + B + 2)$$

$$h(x) > h_2(x) = \frac{B}{(1-x)^2} - (A + B + 2)$$

and we solve for $h_1(x_1) = 0, h_2(x_2) = 0$. This means that

$$x_1 = \sqrt{\frac{A}{A+B+2}}, \quad x_2 = 1 - \sqrt{\frac{B}{A+B+2}}$$

and clearly $h(x_1) > 0, h(x_2) > 0$. Now, if we can show that, for all possible values of A, B in this region $f(x_i) = l_B(x_i) - \max_x l_N(x) > 0$, our proof will be complete.

To make the dependence on A, B more clear we write $f(x_i) = f_i(A, B)$ below

$$f_1(A, B) = \log(1+A) + A \log\left(\sqrt{\frac{A}{A+B+2}}\right) + B \log\left(1 - \sqrt{\frac{A}{A+B+2}}\right) + \frac{1}{2} \log(2\pi) - \frac{1}{2} \log(A+B+2),$$

$$f_2(A, B) = \log(1+A) + A \log\left(1 - \sqrt{\frac{B}{A+B+2}}\right) + B \log\left(\sqrt{\frac{B}{A+B+2}}\right) + \frac{1}{2} \log(2\pi) - \frac{1}{2} \log(A+B+2).$$

We will demonstrate that $\frac{\partial f_i}{\partial B} \leq 0$ for all of the values in our region $A > B > 0$.

$$\begin{aligned}
\frac{\partial f_1}{\partial B} &= -\frac{A}{2(A+B+2)} + \log\left(1 - \sqrt{\frac{A}{A+B+2}}\right) + \frac{B\sqrt{A}}{2(A+B+2)^{3/2}\left(1 - \sqrt{\frac{A}{A+B+2}}\right)} - \frac{1}{2(A+B+2)} \\
&= \frac{1}{2(A+B+2)} \left(\frac{B\sqrt{A}}{\sqrt{A+B+2}\left(1 - \sqrt{\frac{A}{A+B+2}}\right)} - A - 1 \right) + \log\left(1 - \sqrt{\frac{A}{A+B+2}}\right) \\
&= \frac{1}{2(A+B+2)} \left(\frac{B\sqrt{A}}{\sqrt{A+B+2} - \sqrt{A}} - A - 1 \right) + \log\left(1 - \sqrt{\frac{A}{A+B+2}}\right) \\
&\leq \frac{1}{2(A+B+2)} \left(\frac{\sqrt{B}/3}{\sqrt{A+B+2} - \sqrt{A}} - A - 1 \right) - \sqrt{\frac{A}{A+B+2}} \\
&\leq \frac{1}{2(A+B+2)} \left(\frac{1}{3} \sqrt{\frac{B}{B+2}} - A - 1 \right) - \sqrt{\frac{A}{A+B+2}} \\
&\leq -\frac{A}{2(A+B+2)} - \sqrt{\frac{A}{A+B+2}} \leq 0.
\end{aligned}$$

Similarly,

$$\begin{aligned}
\frac{\partial f_2}{\partial B} &= -A \left(\frac{\sqrt{\frac{B}{A+B+2}}}{2B} + \frac{1}{2(A+B+2)} \right) + \log\left(\sqrt{\frac{B}{A+B+2}}\right) + B \left(\frac{A+2}{2B(A+B+2)} \right) - \frac{1}{2(A+B+2)} \\
&= \frac{1}{2(A+B+2)} \left(A+2 - A - 1 - A \sqrt{\frac{A+B+2}{B}} \right) + \log\left(\sqrt{\frac{B}{A+B+2}}\right) \\
&= \frac{1}{2(A+B+2)} \left(1 - A \sqrt{\frac{A+B+2}{B}} \right) + \frac{1}{2} \log\left(\frac{B}{A+B+2}\right).
\end{aligned}$$

Therefore, for any $A \geq 0$ this means that $\frac{\partial^2 f_2}{\partial A \partial B} < 0$. Therefore this expression $\frac{\partial f_2}{\partial B}$ is maximized over A for $A=0$. We can evaluate this expression explicitly:

$$\frac{\partial f_2}{\partial B} \Big|_{A=0} = \frac{1}{2(B+2)} + \frac{1}{2} \log\left(\frac{B}{B+2}\right) \leq \frac{1}{2} \left(\frac{1}{B+2} + \frac{B}{B+2} - 1 \right) \leq 0.$$

This provides a monotonicity result which states that both f_1, f_2 are minimized at at the largest possible $B = \frac{1}{9A}$ for any given A over our region. We will now write $g_i(A) := f_i(A, \frac{1}{9A})$. If we can show that $g_i(A) \geq 0$ for all $A \geq \frac{1}{3}$ and $i=1,2$ we will be done with our proof. We will perform a similar argument to show that g_i is monotone increasing for all $A \geq \frac{1}{3}$.

$$\begin{aligned}
g_1(A) &= \log(1+A) + A \log\left(\sqrt{\frac{A}{A+\frac{1}{9A}+2}}\right) + \frac{1}{9A} \log\left(1 - \sqrt{\frac{A}{A+\frac{1}{9A}+2}}\right) + \frac{1}{2} \log(2\pi) - \frac{1}{2} \log\left(A + \frac{1}{9A} + 2\right) \\
&= \log(1+A) + \frac{A}{2} \log(A) - \frac{1}{2} (1+A) \log\left(A + \frac{1}{9A} + 2\right) + \frac{1}{9A} \log\left(1 - \sqrt{\frac{A}{A+\frac{1}{9A}+2}}\right) + \frac{1}{2} \log(2\pi)
\end{aligned}$$

Note that the function $p(A) = A + \frac{1}{9A}$ is increasing in A for $A \geq \frac{1}{3}$. We can conservatively bound g

from below noting $\frac{1}{9A} \leq 1$ in our region.

$$\begin{aligned} g_1(A) &\geq \log(1+A) + \frac{A}{2} \log(A) - \frac{1}{2}(1+A) \log(A+3) + \frac{1}{9A} \log\left(1 - \sqrt{\frac{A}{A+2}}\right) + \frac{1}{2} \log(2\pi) \\ &\geq \log(1+A) + \frac{A}{2} \log(A) - \frac{1}{2}(1+A) \log(A+3) - \frac{1}{9A} \sqrt{A} + \frac{1}{2} \log(2\pi) =: \tilde{g}_1(A). \end{aligned}$$

We can use calculus to say that:

$$\begin{aligned} \tilde{g}_1'(A) &= \frac{1}{A+1} + \frac{1}{A+3} + \frac{\log(A)}{2} + \frac{1}{18A^{3/2}} - \frac{1}{2} \log(A+3) \\ &\geq \frac{1}{A+1} + \frac{1}{A+3} + \frac{1}{18A^{3/2}} + \frac{1}{2} \log\left(\frac{A}{A+3}\right) \end{aligned}$$

This expression is monotone decreasing in A and with a limit ≥ 0 . Therefore $g_1(A) \geq \tilde{g}_1(A) \geq \tilde{g}_1(1/3)$ for all A . We can explicitly evaluate this numerically and $\tilde{g}_1(1/3) > 0.01$ so we are done. The final piece of this proof involves a similar argument for $g_2(A)$.

$$\begin{aligned} g_2(A) &= \log(1+A) + A \log\left(1 - \sqrt{\frac{\frac{1}{9A}}{A + \frac{1}{9A} + 2}}\right) + \frac{1}{9A} \log\left(\sqrt{\frac{\frac{1}{9A}}{A + \frac{1}{9A} + 2}}\right) + \frac{1}{2} \log(2\pi) - \frac{1}{2} \log\left(A + \frac{1}{9A} + 2\right) \\ &= \log(1+A) + A \log\left(1 - \sqrt{\frac{1}{9A^2 + 18A + 1}}\right) + \frac{1}{2} \left(\frac{1}{9A} \log\left(\frac{1}{9A}\right)\right) - \frac{1}{2} \left(\frac{1}{9A} + 1\right) \log\left(A + \frac{1}{9A} + 2\right) + \frac{1}{2} \log(2\pi) \\ &\geq \log(1+A) + A \left(-\frac{1}{\sqrt{9A^2}}\right) + \frac{1}{2} \left(\frac{1}{9A} \log\left(\frac{1}{9A}\right)\right) - \frac{1}{2} \left(\frac{1}{3} + 1\right) \log\left(A + \frac{1}{3} + 2\right) + \frac{1}{2} \log(2\pi) \\ &\geq \log(1+A) - \frac{1}{3} - \frac{1}{2e} - \frac{2}{3} \log\left(A + \frac{7}{3}\right) + \frac{1}{2} \log(2\pi) =: \tilde{g}_2(A) \end{aligned}$$

Once again we can see that \tilde{g}_2 is monotone increasing

$$\tilde{g}_2'(A) = \frac{1}{1+A} - \frac{2/3}{A+7/3} = \frac{A+5}{(A+1)(3A+7)} \geq 0.$$

We complete the argument by noting $g_2(A) \geq \tilde{g}_2(A) \geq \tilde{g}_2(1/3) > 0.01$. This concludes our proof of the PDF double crossing in region $R3$. \square

The results of Sections 4.2.1, 4.2.2 and 4.2.3 together prove Lemma 3. By proposition 1, Lemmas 2 and 3 together complete the proof of Theorem 1.

5 Acknowledgements

This work was generously supported by a research grant from Boeing, a Marketing Research Award from Adobe, and Stanford Graduate Fellowships, courtesy of PACCAR.

References

- Charles E Antoniak. Mixtures of dirichlet processes with applications to bayesian nonparametric problems. *The annals of statistics*, pages 1152–1174, 1974.
- Josef Hadar and William R Russell. Rules for ordering uncertain prospects. *The American Economic Review*, pages 25–34, 1969.
- Minjung Kyung, Jeff Gill, and George Casella. Characterizing the variance improvement in linear Dirichlet random effects models. *Statistics & Probability Letters*, 79(22):2343–2350, 2009.
- Ian Osband. *Deep Exploration via Randomized Value Functions*. PhD thesis, Stanford, 2016.
- Ian Osband, Benjamin Van Roy, and Zheng Wen. Generalization and exploration via randomized value functions. *arXiv preprint arXiv:1402.0635*, 2014.
- Walter Rudin. *Principles of mathematical analysis*, volume 3. McGraw-Hill New York, 1964.
- Richard Sutton and Andrew Barto. *Reinforcement Learning: An Introduction*. MIT Press, March 1998.